is a protein chemist's art. As always, capturing this art in heuristic rules and putting it to use with an inference engine is the project's goal.

The inference engine for CRYSALIS is a modification of the SU/X system design described above. The hypothesis formation process must deal with many levels of possibly useful aggregation and abstraction. For example, the map itself can be viewed as consisting of "peaks," or "peaks and valleys," or "skeleton." The protein model has "atoms," "amide planes," "amino acid sidechains," and even massive substructures such as "helices." Protein molecules are so complex that a systematic generation-and-test strategy like DENDRAL's is not feasible. Incremental piecing together of the hypothesis using region-growing methods is necessary.

The CRYSALIS design (alias SU/P) is described in a recent paper by Nii and Feigenbaum (1977).

4 SUMMARY OF CASE STUDIES

Some of the themes presented earlier need no recapitulation, but I wish to revisit three here: generation-and-test; situation => action rules; and explanations.

4.1 Generation and Test

Aircraft come in a wide variety of sizes, shapes, and functional designs and they are applied in very many ways. But almost all that fly do so because of the unifying physical principle of lift by airflow; the others are described by exception. So it is with intelligent agent programs and, the information processing psychologists tell us, with people. One unifying principle of "intelligence" is generation-and-test. No wonder that it has been so thoroughly studied in AI research!

In the case studies, generation is manifested in a variety of forms and processing schemes. There are legal move generators defined formally by a generating algorithm (DENDRAL's graph generating algorithm); or by a logical rule of inference (MYCIN's backward chaining). When legal move generation is not possible or not efficient, there are plausible move generators (as in SU/X and AM). Sometimes generation is interleaved with testing (as in MYCIN, SU/X, and AM). In one case, all generation precedes testing (DENDRAL). One case (META-DENDRAL) is mixed, with some testing taking place during generation, some after.

Test also shows great variety. There are simple tests (MYCIN: "Is the organism aerobic?"; SU/X: "Has a spectral line appeared at position P?") Some tests are complex heuristic evaluations (AM: "Is the new concept 'interesting'?"; MOLGEN: "Will the reaction actually take place?") Sometimes a complex test can involve feedback to modify the object being tested (as in META-DENDRAL).

The evidence from our case studies supports the assertion by Newell and Simon that generation-and-test is a law of our science (Newell and Simon, 1976).

4.2 Situation = > Action rules

Situation => Action rules are used to represent experts' knowledge in all of the case studies. Always the situation part indicates the specific conditions under which the rule is relevant. The action part can be simple (MYCIN: conclude presence of particular organism; DENDRAL: conclude break of particular bond). Or it can be quite complex (MOLGEN: an experiential procedure). The overriding consideration in making design choices is that the rule form chosen be able to represent clearly and directly what the expert wishes to express about the domain. As illustrated, this may necessitate a wide variation in rule syntax and semantics.

From a study of all the projects, a regularity emerges. A salient feature of the Situation => Action rule technique for representing expert's knowledge is the modularity of the knowledge base, with the concomitant flexibility to add or change the knowledge easily as the experts' understanding of the domain changes. Here too one must be pragmatic, not doctrinaire. A technique such as this can not represent modularity of knowledge if that modularity does not exist in the domain. The virtue of this technique is that it serves as a framework for discovering what modularity exists in the domain. Discovery may feed back to cause reformulation of the knowledge toward greater modularity.

Finally, our case studies have shown that strategy knowledge can be captured in rule form. In TEIRESIAS, the metarules capture knowledge of how to deploy domain knowledge; in SU/X, the strategy rules represent the experts' knowledge of "how to analyze" in the domain.

4.3 Explanation

Most of the programs, and all of the more recent ones, make available an explanation capability for the user, be he end-user or system developer. Our focus on end-users in applications domains has forced attention to human engineering issues, in particular making the need for the explanation capability imperative.

The Intelligent Agent viewpoint seems to us to demand that the agent be able to explain its activity; else the question arises of who is in control of the agent's activity. The issue is not academic or philosophical. It is an engineering issue that has arisen in medical and military applications of intelligent agents, and will govern future acceptance of AI work in applications areas. And on the philosophical level one might even argue that there is a moral imperative to provide accurate explanations to end-users whose intuitions about our systems are almost nil.

Finally, the explanation capability is needed as part of the concerted attack on the knowledge acquisition problem. Explanation of the reasoning process is central to the interactive transfer of expertise to the knowledge base, and it is our most powerful tool for the debugging of the knowledge base.

5 EPILOGUE

What we have learned about knowledge engineering goes beyond what is discernible in the behavior of our case study programs. In the next paper of this two-part series, I will raise and discuss many of the general concerns of knowledge engineers, including these:

What constitutes an "application" of AI techniques?

> There is a difference between a serious application and an application-flavored toy problem.

What are some criteria for the judicious selection of an application of AI techniques?

What are some applications areas worthy of serious attention by knowledge engineers?

> For example, applications to science, to signal interpretation, and to human interaction with complex systems.

How to find and fascinate an Expert.

> The background and prior training of the expert.

> The level of commitment that can be elicited.

> Designing systems that "think the way I do."

> Sustaining attention by quick feedback and incremental progress.

> Focusing attention to data and specific problems.

> Providing ways to express uncertainty of expert knowledge.

> The side benefits to the expert of his investment in the knowledge engineering activity.

Gaining consensus among experts about the knowledge of a domain.

> The consensus may be a more valuable outcome of the knowledge engineering effort than the building of the program.

Problems faced by knowledge engineers today:

> The lack of adequate and appropriate computer hardware.

> The difficulty of export of systems to end-users, caused by the lack of properly-sized and -packaged combinations of hardware and software

> The chronic absence of cumulation of AI techniques in the form of software packages that can achieve wide use.

> The shortage of trained knowledge engineers.

> The difficulty of obtaining and sustaining funding for interesting knowledge engineering projects.

6 ACKNOWLEDGMENT

The work reported herein has received long-term support from the Defense Advanced Research Projects Agency. The National Institutes of Health has supported DENDRAL, META-DENDRAL, and the SUMEX-AIM computer facility on which we compute. The National Science Foundation has supported research on CRYSALIS and MOLGEN. The Bureau of Health Sciences Research and Evaluation has supported research on MYCIN. I am grateful to these agencies for their continuing support of our work.

I wish to express my deep admiration and thanks to the faculty, staff and students of the Heuristic Programming Project, and to our collaborators in the various worldly arts, for the creativity and dedication that has made our work exciting and fruitful. My particular thanks for assistance in preparing this manuscript go to Randy Davis, Penny Nii, Reid Smith, and Carolyn Taynai.

7 REFERENCES

General

Feigenbaum, E.A. "Artificial Intelligence Research: What is it? What has it achieved? Where is it going?," invited paper, Symposium on Artificial Intelligence, Canberra, Australia, 1974.

Goldstein, I. and S. Papert, "Artificial Intelligence, Language, and the Study of Knowledge," Cognitive Science, Vol.1, No.1, 1977.

Gregory, R., "On How so Little Information Controls so Much Behavior," Bionics Research Report No. 1, Machine Intelligence Department, University of Edinburgh, 1968.

Newell, A. and H.A. Simon, Human Problem Solving, Prentice-Hall, 1972.

Newell, A. and H.A. Simon, "Computer Science as Empirical Inquiry: Symbols and Search," Com ACM, 19, 3, March, 1976.

DENDRAL and META-DENDRAL

Feigenbaum, E.A., Buchanan, B.G. and J. Lederberg, "On Generality and Problem Solving: a Case Study Using the DENDRAL Program," Machine Intelligence 6, Edinburgh Univ. Press, 1971.

Buchanan, B.G., Duffield, A.M. and A.V. Robertson, "An Application of Artificial Intelligence to the Interpretation of Mass Spectra," Mass Spectrometry Techniques and Applications, G.W.A. Milne, Ed., John Wiley & Sons, Inc., p. 121, 1971.

Michie, D. and B.G. Buchanan, "Current Status of the Heuristic DENDRAL Program for Applying Artificial Intelligence to the Interpretation of Mass Spectra," Computers for Spectroscopy, R.A.G. Carrington, ed., London: Adam Hilger, 1974.

Buchanan, B.G., "Scientific Theory Formation by Computer," Nato Advanced Study Institutes Series, Series E: Applied Science, 14:515, Noordhoff-Leyden, 1976.

Buchanan, B.G., Smith, D.H., White, W.C., Gritter, R.J., Feigenbaum, E.A., Lederberg, J. and C. Djerassi, "Applications of Artificial Intelligence for Chemical Inference XXII. Automatic Rule Formation in Mass Spectrometry by Means of the Meta-DENDRAL Program," Journal of the ACS, 98:6168, 1976.

MYCIN

Shortliffe, E. Computer-based Medical Consultations: MYCIN, New York, Elsevier, 1976.

Davis, R., Buchanan, B.G. and E.H. Shortliffe, "Production Rules as a Representation for a Knowledge-Based Consultation Program," Artificial Intelligence, 8, 1, February, 1977.

Shortliffe, E.H. and B.G. Buchanan, "A Model of Inexact Reasoning in Medicine," Mathematical Biosciences, 23:351, 1975.

TEIRESIAS

Davis, R., "Applications of Meta Level Knowledge to the Construction, Maintenance and Use of Large Knowledge Bases," Memo HPP-76-7, Stanford Computer Science Department, Stanford, CA, 1976.

Davis, R., "Interactive Transfer of Expertise I: Acquisition of New Inference Rules," these Proceedings.

Davis, R. and B.G. Buchanan, "Meta-Level Knowledge: Overview and Applications," these Proceedings.

SU/X

Nii, H.P. and E.A. Feigenbaum, "Rule Based Understanding of Signals," Proceedings of the Conference on Pattern-Directed Inference Systems, 1977 (forthcoming), also Memo HPP-77-7, Stanford Computer Science Department, Stanford, CA, 1977.

AM

Lenat, D., "AM: An Artificial Intelligence Approach to Discovery in Mathematics as Heuristic Search," Memo HPP-76-8, Stanford Computer Science Department, Stanford, CA, 1976.

MOLGEN

Martin, N., Friedland, P., King, J., and M. Stefik, "Knowledge Base Management for Experiment Planning in Molecular Genetics," these Proceedings.

CRYSALIS

Engelmore, R. and H.P. Nii, "A Knowledge-Based System for the Interpretation of Protein X-Ray Crystallographic Data," Memo HPP-77-2, Department of Computer Science, Stanford, CA, 1977.

References

1. Adams, J.B. A probability model of medical reasoning and the MYCIN model. Math. Biosci. 32,177-186 (1976).

2. Anderson, R.H., Gallegos, M., Gillogly, J.J., Greenberg, R., and Villanueva, R. RITA Reference Manual, Report R-1808-ARPA, The Rand Corporation, Santa Monica, CA., September 1977.

3. Bennett J.S., Creary L.G., Engelmore R.E., Melosh R.B., A Knowledge-based Consultant for structural analysis, forthcoming.

4. Bleich, H.L. The computer as a consultant. New Eng. J. Med. 284,141-147 (1971).

5. Blum, Robert L. and Wiederhold, Gio: Inferring Knowledge from Clinical Data Banks Utilizing Techniques from Artificial Intelligence. "Proc. 2nd Annual Symp. on Comp. Applic. in Med. Care," pp. 303-307, IEEE, Washington D.C., Nov. 5-9, 1978.

6. Bobrow D.G., Winograd T., An Overview of KRL, a Knowledge Representation Language, Cognitive Science 1:1 (1977).

7. Bobrow D.G., Winograd T., Experience with KRL-0, One cycle of a knowledge representation language, Proceedings of the 5th International Joint Conference on Artificial Intelligence, Cambridge, Mass. (August 1977).

8. Bonnet A., BAOBAB, A parser for a rule-based system using a semantic grammar, Technical Report HPP-78-10, Heuristic Programming Project, Stanford California (September 1978).

9. Brown, J.S., Steps toward a Theoretic Foundation for Complex, Knowledge-Based CAI. BBN No. 3135.

10. Brown, J.S., Collins, A., and Harris, G.

Artificial Intelligence and Learning Strategies. To appear in Learning Strategies (ed. Harry O'Neil), Academic Press, New York, 1978.

11. Buchanan, Bruce G. and Feigenbaum, Edward A. DENDRAL and Meta-DENDRAL: Their Applications Dimension, Artificial Intelligence, 11:5 (1978).

12. Clancey, W. "The Structure of a Case Method Dialogue", to appear in Int. Jnl. of Man Machine Studies, Fall, 1978.

13. Colby, K.M., Weber, S., and Hilf, F. Artificial paranoia. Artificial Intelligence 2,1-25 (1971).

14. Croft, D.J. Is computerized diagnosis possible? Comput. Biomed. Res. 5,351-367 (1972).

15. Davis, R. Applications Of Meta Level Knowledge To the Construction, Maintenance, And Use Of Large Knowledge Bases. Doctoral dissertation, Stanford University ; Memo HPP-76-7, Stanford Computer Science Department, 1976.

16. Davis, R. and King, J. An overview of production systems. Machine Intelligence 8: Machine Representations of Knowledge (eds. E.W. Elcock and D. Michie), John Wiley, April 1977.

17. de Dombal, F.T., Leaper, D.J., Staniland, J.R., McCann, A.P., Horrocks, J.C. Computer aided diagnosis of acute abdominal pain. Brit. Med. J. II,9-13 (1972).

18. Duda, R. O., Hart, P., Nilsson, N. & Sutherland, G. "Semantic network representations in rule-based inference systems", in Pattern Directed Inference Systems (eds. Waterman and Hayes-Roth), Academic Press,New York, 1978.

19. Engelmore R.S., Nii H.P., A knowledge-based system for the interpretation of protein x-ray crystallographic data, Heuristic Programming Project Memo HPP-77-2 (February 1977).

20. Erman L.D., Lesser V.R., A multi-level organization for problem solving using many, diverse, cooperating sources of

knowledge, in Proceedings of the 4th International Joint Conference on Artificial Intelligence, Tbilsi, Russia (1975).

21. Fagan L.M., Ventilator Manager: A program to provide on-line consultative advice in the intensive care unit, Heuristic Programming Project Memo HPP-78-16 (Working Paper), Computer Science Department, Stanford University (September 1978).

22. Feigenbaum E.A., The art of artificial intelligence: I. Themes and case studies of knowledge engineering, Proceedings of the 5th International Joint Conference on Artificial Intelligence, Cambridge, Mass. (August 1977).

23. Feitelson J., Stefik M., A case study of the reasoning in a genetics experiment, Heuristic Programming Project Report 77-18 (working paper) ,Computer Science Department, Stanford University (April 1977).

24. Friedman, R.B. and Gustafson, D.H. Computers in clinical medicine: a critical review. Comput. Biomed. Res. 10,199-204 (1977).

25. Fries, J. Time-oriented patient records and a computer data-bank. J. Amer. Med. Assoc. 222,1536-1542 (1973).

26. Goldstein, I., Papert, S. Artificial Intelligence, Language, and study of knowledge. Cognitive Science 1:1 (1977).

27. Gorry, G.A. and Barnett, G.O. Experience with a model of sequential diagnosis. Comput. Biomed. Res. 1,490-507 (1968).

28. Gorry, G.A., Kassirer, J.P., Essig, A., and Schwartz, W.B. Decision analysis as the basis for computer-aided management of acute renal failure. Amer. J. Med. 55,473-484 (1973).

29. Gorry, G.A., Silverman, H., and Pauker, S.G. Capturing clinical expertise: a computer program that considers clinical responses to digitalis. Amer. J. Med. 64,452-460 (1978).

30. Green, B.F., Wolf, A.K., Chomsky, C., and Laughery, K. BASEBALL: An automatic question-answerer. In Computers and Thought (eds. E.A. Feigenbaum and J. Feldman), pp. 207-216, McGraw-Hill, San Francisco, 1963.

31. Harless, W.G., Drennon, G.G., Marxer, J.J., Root, J.A., Wilson, L.L., and Miller, G.E. CASE - a natural language computer model. Comput. Biol. Med. 3,227-246 (1973).

32. Hart, P.E. Progress on a computer-based consultant. AI Technical Note 99, Stanford Research Institute, Menlo Park, CA., January 1975.

33. Hayes-Roth F., Lesser V.R., Focus of attention in the HEARSAY-II speech understanding system, Proceedings of the 5th International Joint Conference on Artificial Intelligence, Cambridge, Mass. (August 1977).

34. Heiser J.F., Brooks R.E., Ballard J.P., "Progress Report: A Computerized Psychopharmacology Advisor", Proceedings of the 11th Colegium Internationale NeuroPsychopharmacologicum. Vienna, 1978.

35. Heiser, J.F. and Brooks, R.E. A computerized psychopharmacology advisor. Proceedings of the 4th Annual AIM Workshop, Rutgers University, June 1978.

36. Hoffer, E.P. Experience with the use of computer simulation models in medical education. Comput. Biol. Med. 3,269-279 (1973).

37. Kunz J.C., Fallat R.J., McClung D.H., Osborn J.J., Votteri B.A., Nii H.P., Aikins J.S., Fagan L.M., Feigenbaum E.A., A physiological rule based system for interpreting pulmonary function test results, Heuristic Programming Project Memo HPP-78-19, Stanford University, 1978.

38. Lenat D.B., The ubiquity of discovery, Artificial Intelligence 9:3 (1977).

39. Lowerre B.T., The HARPY speech recognition system, Doctoral thesis, Department of Computer Science, Carnegie-Mellon University (April 1976).

40. Martin N., Friedland P., King J., Stefik M., Knowledge Base Management for Experiment Planning, Proceedings of the 5th International Joint Conference on Artificial Intelligence, Cambridge, Mass. (August 1977).

41. Mesel, E., Wirtshcafter, D.D., Carpenter, J.T., Durant, J.R., Henke, C., and Gray, E.A. Clincial Algorithms for Cancer Chemotherapy — Systems for Community-Based Consultant-Extenders and Oncology Centers. Meth. Inform. Med. 15:3, 168-73 (1976).

42. Minsky M., A framework for representing knowledge, in The psychology of computer vision, (ed. P. Winston), McGraw-Hill, New York (1975).

43. Nii H.P., Feigenbaum E.A., Rule-based understanding of signals in Pattern-Directed Inference Systems (eds. Waterman and Hayes-Roth), Academic Press, New York, 1978.

44. Osborn, J.J., Kunz, J.C., and Fagan, L.M. PUFF/VM: interpretation of physiological measurements in the pulmonary function laboratory and the intensive care unit. Proceedings of the 4th Annual AIM Workshop, Rutgers University, June 1978.

45. Pauker, S.G., Gorry, G.A., Kassirer, J.P., and Schwartz, W.B. Towards the simulation of clinical cognition: taking a present illness by computer. Amer. J. Med. 60,981-996 (1976).

46. Pople, H.E., Myers, J.D., Miller, R.A. DIALOG (INTERNIST): a model of diagnostic logic for internal medicine. Proceedings of the 4th International Joint Conference on Artificial Intelligence, pp. 849-855, Tbilisi, Russia, 1975.

47. Quillian, M.R. Semantic memory. In Semantic Information Processing (ed. M. Minsky), pp. 227-270, M.I.T. Press, Cambridge, MA., 1968.

48. Scott, A.C., Clancey, W.J., Davis, R., and Shortliffe, E.H. Explanation capabilities of knowledge-based production systems. Amer. J. Computational Linguistics, Microfiche 62, 1977.

49. Shortliffe, E.H. and Buchanan, B.G. A model of inexact reasoning in medicine. Math. Biosci. 23,351-379 (1975).

50. Shortliffe, E.H., Davis, R., Axline, S.G., Buchanan, B.G., Green, C.C., and Cohen, S.N. Computer-based consultations in clinical therapeutics: explanation and rule-acquisition capabilities of the MYCIN system. Comput. Biomed. Res. 8,303-320 (1975).

51. Shortliffe, E.H. Computer-Based Medical Consultations: MYCIN. Elsevier/North Holland, New York, 1976.

52. Stefik M., An examination of a frame-structured representation system, Stanford Heuristic Programming Project Memo HPP-78-13 (working paper) (September 1978).

53. Stefik M., Inferring DNA structures from segmentation data, Artificial Intelligence 11 (1978).

54. Van Melle, W. Would you like advice on another horn? MYCIN project internal working paper, Stanford University, Stanford, California, December 1974.

55. Warner, H.R., Toronto, A.F., and Veasy, L.G. Experience with Bayes' theorem for computer diagnosis of congenital heart disease. Anns. N.Y. Acad. Sci. 115,558-567 (1964).

56. Weinberg, A.D. CAI at the Ohio State University College of Medicine. Comput. Biol. Med. 3,299-305 (1973).

57. Weiss, S., Kulikowski, C.A., and Safir, A. Glaucoma consultation by computer. Comput. Biol. Med. 8,25-40 (1978).

58. Weyl, S., Fries, J., Wiederhold, G., and Germano, F. A modular self-describing clinical databank system. Comput. Biomed. Res. 8,279-293 (1975).

59. Woods, W.A. et al. The lunar sciences natural language information system: final report, BBN Report 2378, Bolt, Beranek and Newman, Cambridge, MA., June 1972.

60. Wooster, H. and Lewis, J.F. Distribution of computer-assisted instruction materials in biomedicine through the Lister Hill Center Experimental Network. Comput. Biol. Med. 3,319-323 (1973).

61. Wortman, P.M. Medical diagnosis: an information processing approach. Comput. Biomed. Res. 5,315-328 (1972).

62. Yu, V.L., Buchanan, B.G., Shortliffe, E.H., Wraith, S.M., Davis, R., Scott, A.C., and Cohen, S.N. Evaluating the performance of a computer-based consultant. To appear in Computer Programs in Biomedicine, 1978.

63. Yu, V.L., Fagan, L.M., Wraith, S.M., Clancey, W.J., Scott, A.C., Hannigan, J., Blum, R.L., Buchanan, B.G., and Cohen, S.N. Computer-based consultation in antimicrobial selection - a comparative evaluation by experts. Submitted for publication, September 1978.

The appropriate programmatic and administrative personnel of each institution involved in this grant application are aware of the NIH consortium grant policy and are prepared to establish the necessary inter-institutional agreement(s) consistent with that policy.